



PCT


(10) International Publication Number
WO 03/005156 A2

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.



WO 03/005156 A2



(57) Abstract: A system, method, and computer program product for management of bandwidth, cost control, quality of service, and operational efficiency in a fiber optic, ethernet-based, TDMA communications system. A network topology is described in which an aggregating optical node is used between a central office and end users. This allows a variety of topologies between the optical node and end users. Methods of spectral allocation are also described, allowing for improved service and efficient management of bandwidth. Bandwidth can also be managed in the time domain, through a system of bandwidth requests and grants, and through flexible handling of packets, e.g., fragmentation, prioritization, and the use of judiciously sized maximum data units. The invention also features mechanisms by which video service is enhanced for end users. These mechanisms include buffering of video frames, allowing channel switching at the optical node, flexible allocation of bandwidth between video and non-video requirements.

SYSTEM, METHOD, AND COMPUTER PROGRAM PRODUCT FOR
MANAGING COMMUNICATIONS IN ETHERNET-BASED FIBER
OPTIC TDMA NETWORKS

BACKGROUND OF THE INVENTION

5 Field of the Invention

The invention described herein relates to data networks, and more particularly, to the use of ethernet over a fiber optic network.

Related Art

10 One of the current trends in data networking is the use of fiber optic media. Moreover, use of ethernet technology is a practical choice for such networks, given that ethernet is well understood and can be supported by available components. The application of ethernet fiber technology to relatively long distance access networks creates problems, however. Among the unresolved problems is how to share bandwidth efficiently and cost-effectively among multiple users in such an environment. A reasonable
15 quality of service for all users is also desirable. Hence there is a need for a system, method, and computer program product by which bandwidth can be managed in an ethernet-based fiber access network, and service can be kept affordable and user -friendly to end users.

20 SUMMARY OF THE INVENTION

The invention described herein addresses management of bandwidth, cost control, quality of service, and operational efficiency in a fiber optic, ethernet-based, TDMA communications system. A network topology is described in which an aggregating optical node is used between a central
25 office and end users. This allows a variety of topologies between the optical node and end users. Methods of spectral allocation are also described, allowing for improved service and efficient management of bandwidth. Bandwidth can also be managed in the time domain, through a system of bandwidth requests and grants, and through flexible handling of packets, e.g.,

- 2 -

fragmentation, prioritization, and the use of judiciously sized maximum data units. The invention also features mechanisms by which video service is enhanced for end users. These mechanisms include buffering of video frames, allowing channel switching at the optical node, flexible allocation of bandwidth between video and non-video requirements.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a block diagram illustrating a gigabit media-independent interface (GMII) incorporated into a media access controller (MAC).

FIG. 2 illustrates a fiber optic network using an active passive optical network (PON) architecture.

FIG. 3 is a block diagram illustrating the use of an optical node to accommodate a variety of communications topologies.

FIG. 4 is a block diagram illustrating an optical node.

FIG. 5 illustrates the use of an adaptive equalizer to reduce noise in a light source.

FIG. 6 is a flowchart illustrating the process of timebase synchronization where a remote maintains synchronization by locking on to a phase-locked transmit bit stream.

FIG. 7 is a flowchart illustrating the use of a synchronization byte in a Moving Pictures Expert Group (MPEG) frame to synchronize a remote device.

FIG. 8 is a flowchart illustrating the use of a synchronization byte in an variable length packet to synchronize a remote device.

FIG. 9 is a flowchart illustrating the gating of upstream transmissions according to a grant of a headend.

FIG. 10A and 10B illustrate the relationship between PON arbitration and the 802.3 protocol.

FIG. 11 illustrates the use of different wavelengths to carry downstream video data, downstream non-video data, and upstream data.

FIG. 12 illustrates the concept of spectral slicing.

- 3 -

FIG. 13 illustrates hybridization of point-to-point and broadcast architectures.

FIG. 14 illustrates subcarrier multiplexing, wherein each user has its own subcarrier.

5 FIG. 15 is a flowchart illustrating the use of a maximum data unit (MDU) in bandwidth allocation.

FIG. 16 is a flowchart illustrating flexible packet fragmentation based on available bandwidth.

10 FIG. 17 is a flowchart illustrating reallocation of video bandwidth for non-video data.

FIG. 18 is a flowchart illustrating the bandwidth request and grant process.

FIG. 19 is a flowchart illustrating packet transmission based on priority and bandwidth availability.

15 FIG. 20 is a block diagram illustrating the buffering of MPEG frames at an optical node (ON).

FIG. 21 is a flow chart illustrating proactive video streaming.

DETAILED DESCRIPTION OF THE INVENTION

20 A preferred embodiment of the present invention is now described with reference to the figures, where like reference numbers indicate identical or functionally similar elements. Also in the figures, the left-most digit of each reference number corresponds to the figure in which the reference number is first used. While specific configurations and arrangements are discussed, it should be understood that this is done for illustrative purposes only. A person skilled in the relevant art will recognize that other configurations and arrangements can be used without departing from the spirit and scope of the invention. It will be apparent to a person skilled in the relevant art that this invention can also be employed in a variety of other devices and applications.

25

- 4 -

I. Overview

In the invention described herein, the use of a time division multiple access (TDMA) architecture allows the sharing of bandwidth among multiple users. The Data Over Cable System Interface Specification (DOCSIS) describes a process through which bandwidth management and other requirements can be achieved in a TDMA setting. The present invention provides means for addressing management of bandwidth, cost control, quality of service, and operational efficiency in a fiber optic, ethernet-based, TDMA communications system.

In general, DOCSIS can be used in communication systems that include a set of remote communications devices connected to a headend device, such that the headend is responsible for the management of communications both to and from the remotes. The headend is responsible for the distribution of information content to the remotes (the so-called "downstream" direction); in addition, the headend is responsible for management of communications in the other direction, from the remotes to the headend (the "upstream" direction). Generally, in addition to sending content to remotes, the headend issues downstream map messages that instruct each remote as to when it can transmit upstream, and what kind of information it can send. In effect, the upstream bandwidth is controlled and allocated by the headend. Any given remote can transmit upstream only after requesting bandwidth and receiving a grant of the bandwidth from the headend. In a time division multiple access (TDMA) environment, bandwidth corresponds to one or more intervals of time. Moreover, the upstream can be organized into a number of channels, with several remotes assigned to each channel. This arrangement allows the headend to manage each upstream communications channel. In this manner, upstream communications are managed so as to maintain order and efficiency and, consequently, an adequate level of service.

In the realm of cable communications, DOCSIS specifies the requirements for interactions between a cable headend and associated remote

- 5 -

cable modems. A cable headend is also known as a cable modem termination system (CMTS). DOCSIS consists of a group of specifications that cover operations support systems, management, and data interfaces, as well as network layer, data link layer, and physical layer transport. Note that DOCSIS does not specify an application layer. The DOCSIS specification includes extensive media access layer and physical (PHY) layer upstream parameter control for robustness and adaptability. DOCSIS also provides link layer security with authentication. This prevents theft of service and some assurance of traffic integrity.

The current version of DOCSIS (DOCSIS 1.1) uses a request/grant mechanism for allowing remote devices (such as cable modems) to access upstream bandwidth. DOCSIS 1.1 also allows the provision of different services to different parties who may be tied to a single modem. With respect to the processing of packets, DOCSIS 1.1 allows segmentation of large packets, which simplifies bandwidth allocation. DOCSIS 1.1 also allows for the combining of multiple small packets to increase throughput as necessary. Security features are present through the specification of 56-bit Data Encryption Standard (DES) encryption and decryption, to secure the privacy of a connection. DOCSIS 1.1 also provides for payload header suppression, whereby repetitive ethernet/IP header information can be suppressed for improved bandwidth utilization. DOCSIS 1.1 also supports dynamic channel change. Either or both of the downstream and upstream channels can be changed on the fly. This allows for load balancing of channels, which can improve robustness.

While the present invention is described herein in the context of DOCSIS, it should be understood that the systems and methods discussed below are also applicable in other contexts as well. Generally, these systems and methods are applicable to any fiber access system.

Note that in the discussion below, techniques are organized generally according to their benefit, i.e., cost control, bandwidth management, user-friendliness, and operational efficiency. This does not necessarily represent a

- 6 -

limitation on the utility or scope of any of the techniques. A bandwidth management technique may, for example, have benefits with respect to a system's operational efficiency or user-friendliness. The categorization below should therefore not be viewed as any sort of limitation of applicability.

5

II. Cost Control

A. Hardware architecture

1. Detection of reference point at GMII

10

This aspect of the invention allows the use of existing, commercially available ethernet physical layer devices in a fiber optic TDMA network operating under DOCSIS. DOCSIS includes a process whereby a headend and associated remote devices become synchronized so that they all share the same sense of time with respect to upstream communications. This synchronization is known in DOCSIS as ranging. Ranging requires that remotes each send a burst of information at a time prescribed by the headend. The headend must then detect whether the burst arrived too soon or too late, relative to the prescribed arrival time. Typically, a specific reference point in the preamble of the burst is used to gauge the burst's arrival. When this point is detected, the burst is considered to have arrived. The start of frame delimiter (SFD) byte in a DOCSIS burst can be used for this purpose.

15

20

Commercially available ethernet physical layer devices, however, do not typically have the functionality that allows detection of a burst's reference point. In this invention, the reference point is observed not at the PHY device per se, but rather at the interface of the physical layer device to the media access controller. Because the delay through an ethernet physical layer device is nearly constant, however, it is not necessary for this device to detect the reference point.

25

30

In an embodiment of the invention, illustrated in FIG. 1, the detection takes place at a gigabit media-independent interface (GMII) 120 of the media access controller (MAC) 110. This can be implemented by hardware in MAC

- 7 -

110 that monitors incoming data from PHY device 130, to detect an SFD byte. The timing reference in the MAC 110 can be used to determine the time at which the SFD is observed, relative to the expected time of the SFD's arrival. The GMII and MAC may be placed at an optical node and/or at a headend device. Note that, in general, a GMII can be used to detect any field having a known relationship to the start of a frame, not just an SFD per se, for purposes of detecting a ranging offset.

2. Active PON

Another issue in the use of a fiber access network is the management and allocation of costs in topologies involving relatively long distances (e.g., 20 kilometers or more). One way to address this is to use an active architecture, instead of a passive optical network (PON) approach.

Traditionally, in a PON, transmission of information between a central office and end users, e.g., in their homes, is done through a passive splitter. While this is a workable architecture for relatively short distances, longer distances, up to 20 kilometers and beyond, are problematic. Longer distances require more powerful (and more expensive) light sources.

Instead, an optical node acting as an aggregation device can be used to handle transmissions over longer distances. An embodiment of such a system is shown in FIG. 2. An upstream feeder channel 205 and a downstream channel 210 are shown, each operating on a different wavelength. The downstream feeder channel 210 connects a central office (CO) 215 and an optical node (ON) 220. ON 220 serves as an aggregation device. A distribution system 225 extends from ON 220 to users 230. Thus, a single expensive laser connects ON 220 and CO 215. Less powerful and, consequently, cheaper lasers in distribution system 225 connect end user devices 230 (e.g., modems) to ON 220.

The relatively high cost of the long distance laser can now be shared among users 230. The distribution path from ON 220 to users 230 is relatively

- 8 -

cheap, since less power is required for shorter distances. The costs of using a remote device can therefore be lowered by this approach.

Moreover, in an embodiment of the invention, an optical node can accommodate multiple topologies on the user side. This is illustrated in FIG. 3. A hub 305 sends high bandwidth data, such as digital video, to an ON 310 via a feeder link 315. An end user in a home 320 can receive the data from ON 310 via a fast ethernet point-to-point (P2P) connection 325. Other homes, such as homes 330a through 330n, receive data from ON 310 through a shared connection 335.

An embodiment of an ON is shown in FIG. 4. ON 400 includes a gigabit ethernet interface 405 through which ON 400 connects with a hub. Interface 405 is connected to a switch 410. Switch 410 serves to aggregate traffic headed upstream towards a hub, and can be implemented as a multiplexer/demultiplexer. Switch 410 is connected to one or more PON controllers 415a through 415m, which arbitrate access and provide link control with respect to end users. Switch 410 and PON controllers 415 provide quality of service functions as well, such as the control of data flow based on prioritization or based on other categorizations of traffic. Each PON controller is connected to one or more PHY devices 420a through 420p. Each PHY device is then connected through a physical communications medium to an end user device (not shown).

Moreover, in an embodiment of the invention, ON 400 has other interfaces (not shown), to support different kinds of traffic, such as voice, and/or to support circuit emulation.

B. Operational efficiency

1. Cancellation of Laser Humming

In any optical access system, the light sources (e.g., lasers) may not operate continually. Rather, they can cycle as necessary between a powered operational state and an idle state. In the latter, a laser is not completely

- 9 -

powered down. The laser emits light at a low level during idle, and is said to "hum." Humming adds noise, affecting the signal-to-noise ratio (SNR) of other signals in the system.

5 This noise can be ameliorated by using an adaptive equalizer. As is known in the art, an adaptive equalizer can be used to cancel noise on a communications channel. Such an equalizer can cancel the humming of a laser during idle, thereby improving the SNR of information-bearing signals. In an embodiment of the invention, an adaptive equalizer is used as illustrated in FIG. 5. A receiver 500 receives light 505 from a light source, such as a
10 laser. Light 505 is received at an optical detector 510. The output of optical detector is fed to adaptive equalizer 530. Equalizer output 530 produces an equalized output 550.

2. Spectral Slicing

15 Spectral slicing is a technique by which multiple users can use different frequency bands of the same broadband laser source for communication. This is illustrated generally in FIG. 12. Users transmit using different frequencies 1205. These frequencies represent slices of a broadband transmission 1210. This technique enables the implementation of point-to-point links in a point-to-multipoint topology. Since each subscriber uses a
20 different frequency band, subscribers do not interfere with each other.

In such a system there is a tradeoff between the number of subscribers, the bandwidth of the filters required, and the transmit power from each subscriber. There is also a cost tradeoff based on the Q of the filters required.
25 A technique is used in which each subscriber unit has multiple light emitting diodes (LEDs). In an embodiment of the invention, four LEDs are used by subscribers, red, blue, green and yellow. A subscriber can use any one of them for communication with the hub/ON. Since LEDs are very cheap, they will not add significant cost to the subscriber unit. Each unit uses lower Q
30 filtering (representing lower cost) and, as a result, gets to use higher transmit

- 10 -

power. The receiver in the ON can split the four different wavelengths using devices like Briggs grating and can demultiplex different subscribers in each wavelength using filters. This enhances the efficiency of bandwidth usage and increases the number of subscribers per port at hub/ON. In addition, this
5 reduces the cost of the overall system.

The split ratio of the PONs can be increased by using signal processing techniques. By using forward error correction (FEC), coding gains on the order of 3-6 dB can be achieved. This can easily double or quadruple the number of subscribers on a single PON. A further improvement of 3dB can be
10 achieved using adaptive equalization, which can double the subscribers. Since these signal processing techniques can be adding at very little additional cost, the overall cost of the system per subscriber drops significantly.

3. Allocation of Functionality Between Hub and Optical Node

15

An ON, like any other communications component, has limits as to the functionality that it can incorporate. Factors such as chip size and power dissipation must be considered during system design. DOCSIS, however, requires certain functionality at a headend. This includes timing and
20 sequencing functions, such as ranging. DOCSIS also requires bandwidth allocation processing, such as the generation of map messages. It also requires subscriber service functions, such as authentication and billing.

Because all this functionality can be difficult to put in a single component, a better approach is the dispersal of the functionality. In the
25 context of an optical network such as that of FIG. 3, some functionality, such as the subscriber service functions, can be placed in the hub 305. Other functions, such as timing and bandwidth allocation functions, can be placed in the ON 310. This reduces the processing burden on any single component, with no loss in overall system capability.

30

- 11 -

C. Bandwidth Management

1. Timebase Synchronization

TDMA systems require the maintenance of a time base which is used to determine time slot boundaries, communicate the time base to all the equipment in the system, and chronologically lock equipment to the time base. The current state of the art is exemplified by the DOCSIS specification. In such systems, the headend generates a time base in the form of a time stamp counter driven by a very precise reference oscillator. The headend communicates the time base to one or more remote devices (e.g., cable modems) via periodic synchronization messages. These messages contain the current time stamp counter value. There are several problems with such a system. Among them, time stamps must be sent relatively often, and the time, as maintained at a remote, can drift slowly so that it can move several counts away from the headend's count. Recovering from such a variation can take a long period of time.

One method of maintaining synchronization in a TDMA system is the use of synchronous rate locking to keep the rates of the time stamp counts at the headend and each remote device locked to each other. One embodiment uses ethernet PHY devices at both ends. The transmitting PHY at the headend can be viewed as the master. This method is illustrated in FIG. 6. The method starts at step 605. In step 610, the transmitting PHY's transmit bit rate (i.e., the symbol rate for optical PHY) is phase locked to the clock used to generate the headend time stamp counter. In step 620, transmission begins. In step 625, the receiving PHY device at the remote locks on to the bit rate of the incoming data stream. From this the remote's clock is derived locally, which drives the remote's local copy of the time stamp counter (step 630). The method concludes at step 640. Using this method, synchronization messages need not be sent often since they are only used to initialize the counter of a remote when it joins the network and to periodically check the counter against the current value. To initialize, the remote simply loads the first time stamp it

- 12 -

receives into its local register. Techniques such as block coding or scrambling can be used to control clock jitter. Block coding has the advantage of maintaining DC balance and can also maintain the required number of bit transitions. Scrambling techniques can have much less overhead.

5 Other techniques by which time base synchronization can be maintained in a TDMA optical system include an increased frequency of synchronization messages to deal with jitter. This imposes a requirement of the time base generator at the headend to be accurate within 100 picoseconds. This also requires the remote to have tight control on its jitter.

10 Another option is to use physical layer in-band synchronization using MPEG framing. An MPEG frame has a synchronization byte at the start of the frame. This byte has a specific predetermined synchronization pattern. A remote latches to the periodic synchronization byte to synchronize to the downstream rate. This process is illustrated in FIG. 7, according to an embodiment of the invention. The process begins with step 710. In step 720,
15 a counter is initialized. This counter is used to count the number of times, in succession, that the synchronization pattern is successfully found. In step 730, the remote device searches for the synchronization pattern in incoming traffic. If the pattern is not found, as determined in step 740, searching continues at
20 step 730. If the pattern is found, processing continues at step 750, where the counter is incremented. In step 760, a determination is made as to whether the counter has reached a threshold value. If not, processing continues at step 770. Here, the synchronization pattern is sought at a subsequent point in the traffic, a predetermined number of bytes later. For fixed-length MPEG frames, the
25 pattern is sought 188 bytes later. If, in step 780, the synchronization pattern is found, the counter is incremented in step 750, and the process repeats from this point. If no synchronization pattern is found in step 780, the counter is re-initialized in step 720, and the entire process restarts. If, in step 760, the threshold is reached, this indicates that a sufficient number of synchronization
30 patterns have been found in consecutive attempts, and synchronization is attained (step 790).

- 13 -

A similar technique can be implemented with variable length packets with the synchronization byte followed by a pointer to the next synchronization byte. This is illustrated in FIG. 8. The process begins with step 810. In step 820, a counter is initialized. Again, this counter is used to count the number of times, in succession, that the synchronization pattern is successfully found. In step 830, the remote device searches for the synchronization pattern in incoming traffic. If the pattern is not found, as determined in step 840, searching continues at step 830. If the pattern is found, processing continues at step 850, where the counter is incremented. In step 860, a determination is made as to whether the counter has reached a threshold value. If not, processing continues at step 865. Here, a pointer is read, where the pointer is found after the last synchronization pattern. The pointer indicates the location, in the incoming traffic, of the next synchronization pattern. In step 870, the next synchronization pattern is sought at the indicated point in the traffic. If, in step 880, the synchronization pattern is found, the counter is incremented in step 850, and the process repeats from this point. If no synchronization pattern is found in step 880, the counter is re-initialized in step 820, and the entire process restarts. If, in step 860, the threshold is reached, this indicates that a sufficient number of synchronization patterns have been found in consecutive attempts, and synchronization is attained (step 890).

2. Wavelength Allocation, Video and Data

In the communications systems described herein, bandwidth limitations can be problematic. Given one gigabit per second of downstream bandwidth, for example, 600 megabits could be required for digital video, leaving only 400 megabits for other data traffic. Typically, video and data signals share the bandwidth through a multiplexing arrangement.

An alternative is to allocate different wavelengths to different requirements. For example, one wavelength could be allocated to downstream digital video, while another wavelength would be allocated to downstream

- 14 -

non-video data. A third could be allocated to upstream data. This increases the available bandwidth for each requirement, and represents a way to upgrade a traditional PON architecture in light of the need for greater capacity. This is illustrated in FIG. 11. Here, a central office 1110 is in communication with optical node 1120. Downstream digital video is carried on channel 1130, operating at a wavelength λ_1 . Downstream data (non-video) is carried on channel 1140, operating on a wavelength λ_2 . Channel 1150 is used for upstream communications on a wavelength λ_3 . Such an arrangement serves to increase bandwidth between central office 1110 and a set of users 1160.

Note that a video transmission from central office 1110 can be a broadcast, so that multiple optical nodes may receive the λ_1 transmission. Allocation of wavelengths for downstream non-video transmissions (in FIG. 11, wavelength λ_2) and upstream transmissions (wavelength λ_3), however, is done per optical node.

3. Hybrid PON: Broadcast Downstream, Point-to-Point Upstream

Another architectural solution to the bandwidth constraint issue is to hybridize broadcast and point-to-point concepts. This is illustrated in FIG. 13. A central office (CO) 1305 broadcasts downstream to all end users, including a remote device 1310, shown here as customer premises equipment (CPE). The broadcast takes place using a single wavelength, λ_d , and passes through a series of couplers, including couplers 1315 and 1320.

Upstream transmissions take place over multiple wavelengths, one per user, shown here as λ_{u1} through λ_{un} . Hence the upstream is a point-to-point architecture using wavelength division multiplexing (WDM).

Here, the need for a high-powered laser is limited to the CO 1305, as is the need for wavelength detection functionality. Remote devices, such as CPE 1310, require a high bandwidth receiver (e.g., gigabit), but can operate with a lower bandwidth transmitter (e.g., 10/100 megabit).

- 15 -

4. Subcarrier Multiplexing

Where upstream bandwidth is problematic, each user can be assigned his or her own frequency, such that all user frequencies are associated with a single narrowly defined wavelength range. Frequencies can be offset, for example, by 100 MHz in an embodiment of the invention. This allows autonomous communication for each user, without interference. This is illustrated in FIG. 14. Here, a central office 1410 is in communication with users 1461, 1462, and 1463. Each of these three respective users can transmit to central office 1410 using subcarriers 1451, 1452, and 1453, respectively.

5. PON Protocol Architecture: Reservation Ethernet

One possibility for a protocol architecture for PON is the use of a reservation ethernet approach. Here, a gating transmission is used, based on a request grant mechanism on top of ethernet. This approach is illustrated in FIG. 9, beginning with step 905. As in a DOCSIS-like protocol, the ethernet switch generates a map message or grant in step 910, to indicate to the remotes when to transmit. A remote receives the grants and determines, in step 915, whether transmission can take place. If not, the remote uses a gating mechanism in step 925 to hold the ethernet transmission since the remote is not allowed to transmit. Otherwise, in step 920, the remote sends the transmission during the granted periods. The process concludes at step 930. Hence this mechanism arbitrates access between remotes, but keeps the underlying ethernet framing transmission. The additional functionality required at the ethernet switch is the gathering and scheduling of requests, creation of acknowledgment responses, creation of map messages, and transmission of the messages downstream. The functionality required at the remote is the reception and interpretation of the map messages, creation of requests, and the gating mechanism to open or block the ethernet transmission. The control messages (such as grants and acknowledgments) generated by the

- 16 -

ethernet switch can be specified as new ethernet control messages. To be fully ethernet compliant and avoid fragmentation of frames, minimum grant size can be of a size to fit a payload equal to a minimum ethernet frame size (64 bytes). The message-carrying requests can also be defined to be of a size
5 equal to this minimum frame size. Since request messages are small, this message can be specified to allow the carrying of more than one request at a time.

This gating mechanism based on grant messages from the CO (via, e.g., an OLT) to the ON (e.g., ONU) defines a basic communication between
10 the two. Once the CO recognizes the ON, this mechanism assigns a minimum amount of bandwidth to each ON. Additionally, the ON can request more bandwidth as needed. Hence, this mechanism has a contention approach only when the ON is recognized in the system. After this, the access of recognized ONs is contention free. The amount of bandwidth assigned to a recognized
15 ON can be set at a fixed level when the ON is recognized. This amount can be different for each ON depending on the service agreement given to the ON. In addition, the ability to modify this agreement can be defined in order to allow modification of services more dynamically than just during registration time.

Another option for the PON protocol architecture is the use of
20 reservation aloha (request grant mechanism) as the underlying transmission mechanism. The protocol can be defined as a simple version of DOCSIS with the minimum features in it. For example, fragmentation, payload header suppression, and downstream MPEG transport can be eliminated.

Possible relationships between PON arbitration and the 802.3 protocol
25 are illustrated in FIG. 10A and 10B. FIG. 10A illustrates the reservation ethernet case, discussed above, according to one embodiment of the invention. Here a PON arbitration process 1005 provides transmission control inputs 1010 to the 802.3 protocol. These inputs can be translated to ethernet frames. And hence, the definition of this protocol just reduces to define the new frame
30 types to carry this additional arbitration information. Alternatively, the 802.3 protocol could be modified to incorporate the PON arbitration in a lower layer

- 17 -

of the protocol stack as shown in FIG. 10B. A DOCSIS approach would define the PON arbitration as an additional encapsulation mechanism. Another approach is to consider the PON arbitration as physical layer signaling (such as invalid PCS codes in ethernet).

5 FIG. 10B illustrates the protocol relationships in the reservation aloha case, discussed above. Here, PON arbitration 1005 is below the 802.3 protocol, which in turn is below IP layer 1020. Therefore, FIG. 10B illustrates an architectural definition of a DOCSIS PON (DPON), wherein an additional header or protocol is provided for PON arbitration. On the other hand, FIG.
10 10A illustrates an architectural definition for an ethernet-based PON (EPON), wherein the ethernet protocol is extended to provide PON arbitration.

6. Limiting PDU Size and Controlling Fragmentation

15 Under DOCSIS, when a remote receives a grant, it transmits packets in its queue. In DOCSIS there is a one-to-one mapping between the grant and the request. Hence the transmitted packets correspond to the granted bandwidth except for a small amount of bandwidth due to the minislot-to-grant granularity.

20 This one-to-one mapping is not available if more smart mechanisms are available in the system. For example, the headend may generate additional unsolicited grants. If a flexible use of grants is implemented, any "flow" can use any grant independently of which flow generated the request. In this case, the granted bandwidth can be filled up with packets until no more packets fit.
25 At the end of the burst it will leave a space that may not fit the next packet to be transmitted.

 There are generally two options. First, the space can be left unused. This is inefficient. Second, the next packet could be fragmented. Therefore, a system with no fragmentation may be inefficient if the burst lengths are not
30 large enough. On average there is a waste of half of a packet of average size,

- 18 -

per burst. Depending on the burst and average packet sizes, this can be a significant waste.

An alternative to the customary fragmentation approach is to coordinate the packetization at a higher protocol level and specify the sizes of unsolicited grants. In other words a maximum data unit (MDU) can be imposed to break the transmitted data into units that can be better handled in the system. This process is illustrated in FIG. 15. The process begins at step 1510. In step 1520, an MDU is determined to have a size equal to M bytes. In step 1530, the amount of bandwidth to be granted is defined, as an integer multiple of the MDU size. In step 1540, a grant is issued, specifying bandwidth equal to K times the MDU size. The process concludes at step 1550. By carefully choosing the size of unsolicited grants and the MDU size (e.g., defining the size of the grant to be a multiple of the MDU), wasted bandwidth can be minimized. In an embodiment of the invention, different remotes in the system (and even different flows) can operate with different MDU values.

Another alternative is to perform fragmentation in a more flexible manner, depending on the bandwidth available. Again, what is typically done at the media access layer is now done at the IP layer. In this alternative, an arriving grant is examined to identify its size. An IP packet is then fragmented so as to fit the grant, and the IP header is modified as necessary. This is illustrated in FIG. 16. The process begins at step 1610. In step 1620, a remote device receives a grant of bandwidth. In step 1630, the remote device determines the size of the granted bandwidth. In step 1640, a determination is made as to the amount of granted bandwidth that can be used by whole IP packets. This step determines the number of IP packets that can be contained in the granted bandwidth, and calculates the amount of bandwidth that is consumed thereby. In step 1650, the remaining bandwidth is determined. In step 1660, the next IP packet is fragmented so as to use the remaining bandwidth. The process concludes at step 1670. In an embodiment of the

invention, this adaptive process is implemented in hardware, and can be performed in real time.

7. Using Video Bandwidth for Data

5 As is apparent from the above discussion, the proper allocation of bandwidth is required to service a set of users that have a variety of needs. Ideally, allocation of bandwidth is flexible to allow servicing of different needs as they arise.

10 Digital video represents a large amount of data transmitted in a continual stream, and therefore requires significant bandwidth. Accordingly, digital transmissions are generally allocated large amounts of bandwidth by default. But, under some circumstances, requirements for non-video data may be great enough to exceed the default allocations for such data. In this case, bandwidth can be taken from transmissions such as video, and reallocated to
15 data channels that require more bandwidth. This requires monitoring of the demand for non-video data. If a predefined demand threshold is exceeded for non-video data, reallocation takes place. If and when such demand returns to a predefined lower level, the system can return to its default bandwidth allocations. In an embodiment of the invention, the reallocation of video
20 bandwidth for other data transmissions can also depend on whether the demand for video is sufficiently low. This process is illustrated in FIG. 17. The process begins with step 1710. In step 1720, a determination is made as to whether the current requirement for non-video bandwidth exceeds a default value. If not, the normal default allocation for non-video bandwidth is used in
25 step 1730. If the requirement for non-video bandwidth exceeds the default value, however, the process continues at step 1740. Here, a determination is made as to whether the demand for video bandwidth is sufficiently low so as to permit reallocation of video bandwidth to non-video data. If the demand for video is not sufficiently low, then the process continues at step 1730, and the
30 normal allocation of non-video data bandwidth is used. If, however, the demand for video is sufficiently low to allow reallocation, then the process

- 20 -

continues at step 1750. Here, video bandwidth is reallocated for non-video data. The process then returns to step 1720 for continued monitoring of the requirement for non-video bandwidth.

5 8. Allocating Bandwidth With Requests and Grants

 In an embodiment of the invention, bandwidth can be allocated flexibly by using a request/grant mechanism. Such an arrangement is currently defined in the DOCSIS 1.1 standard, but the concept can be adapted to a non-DOCSIS system as well. In such an arrangement, a remote device
10 seeking to transmit does so after requesting bandwidth from a central authority, such as a headend or similar module. If bandwidth is available, a grant is made by the central authority to the remote, specifying the bandwidth to be used by the remote (e.g., a specific time interval). This process is illustrated generally in FIG. 18. The process begins with step 1810. In step
15 1820, a remote device requests bandwidth from the headend. In step 1830, a determination is made by the headend as to whether bandwidth is available. If not, a subsequent request for bandwidth can be made in step 1820. If bandwidth is available, then the process continues at step 1840, where the request for bandwidth is granted, and the amount of bandwidth is specified in
20 the grant. The process concludes at step 1850.

 In an embodiment of the invention, some or all grants can be unsolicited. During registration, bandwidth is allocated according to a fixed assignment policy. As such, the headend can make unsolicited, fixed bandwidth allocations based on state for each remote device. When additional
25 remote devices register, the headend assigns the bandwidth allocation based on availability. In embodiments, the headend dynamically adjusts the bandwidth allocations as the system conditions change, such as remote devices terminating or initiating sessions. The CO keeps the state of the ON bandwidth needs based on the established sessions. In embodiments, the headend
30 dynamically adjusts the bandwidth allocations in response to requests. The

- 21 -

adjustment can be in accordance with established dynamic service level agreements with the remote devices.

Contention among remotes for granted bandwidth can be resolved through a priority system or other mechanism. Note that in a TDMA context,
5 the remote and headend must share the same sense of time. This allows a remote's sense of a granted timeslot (starting and ending points) to match that of the headend. Hence a synchronization process may be required prior to any actual request / grant processing.

10 9. Re-prioritization of Packets to Use Available Bandwidth

In some communications systems, a priority system is in place to resolve contention for available bandwidth. A packet having the highest priority will generally be allowed to use the bandwidth, instead of other lower
15 priority packets that may need to be sent. In some situations, however, this can be an inefficient arrangement. The highest priority packet may be larger than the amount of available bandwidth. The priority logic dictates that only the highest priority packet can be sent, yet this packet cannot be sent because of its size. In this case, the available bandwidth may go wasted.

20 To address this, an exception can be made to the normal priority rules. If a lower priority packet will fit the available bandwidth, this packet will be sent instead of the higher priority packet, rather than wasting the bandwidth. In an embodiment of the invention, the packet to be sent can be identified by choosing the highest priority packet among those that fit the available
25 bandwidth. This is illustrated in FIG. 19. The process begins with step 1910. In step 1920, a remote device receives a bandwidth allocation. In step 1930, the remote device identifies the highest-priority packet among the packets that need to be sent. In step 1940, a determination is made as to whether the highest priority packet fits the allocated bandwidth. If not, then in step 1950,
30 the highest-priority packet is withdrawn from consideration, since it would not fit the allocated bandwidth. The process would then continue at step 1930,

- 22 -

where, among the remaining packets, the highest-priority packet is identified. If, in step 1940, the highest-priority packet fits the allocated bandwidth, then the process continues at step 1960. Here, the packet is sent. The process concludes at step 1970. Alternatively, if efficient bandwidth usage is important, the system can choose the largest packet that will fit. Alternatively, some combination of best fit and highest priority can be used to determine the packet to be sent.

III. User Services

A. Video Switching at Optical Node

Users typically desire the ability to readily control what information they access. In the context of downstream digital video, this includes the ability to select a channel for viewing. Current architectures provide for switching at a hub, such as hub 305 of FIG. 3. In response to a user command, hub 305 performs the requested switching and forwards the appropriate transmission to ON 310, and ultimately to the user.

This creates latency in system response to the user's commands, however, given that the command must go all the way to the hub 305, which must then react. Alternatively, the link 315 carries broadcast video of all transmissions to ON 310. Switching is then performed there, instead of at hub 305. While this requires greater bandwidth between hub 305 and ON 310, the latency of the response to user input is reduced. Moreover, this switching function can also be performed at a central office if, for example, the system does not include an optical node.

B. MPEG Buffering at Optical Node

When MPEG-formatted video is transmitted, a sequence of individual frames is organized into a "group of pictures" (GOP). A GOP begins with an I frame, and is followed by B frames (or T frames, depending on the method of coding). Generally, if a user switches to a transmission at a time when a GOP

- 23 -

has already started, i.e., after the I frame, the entire GOP associated with that I frame is inaccessible.

This can be remedied if GOPs are buffered. This is illustrated in FIG. 20. Here, headend 2010 transmits a GOP 2020. GOP 2020 is held in
5 buffer 2040, located in optical node 2030. This makes each frame of GOP 2020 available to user 2050. A user switching to a transmission in mid-GOP can then access a full GOP, starting with its I frame. In an embodiment of the invention, the GOP is buffered at the ON in a circular buffer. When a user switches to a video transmission, he or she has access to all of the current
10 GOP, since all the GOP's frames up to this point, starting with the GOP's I frame, are available.

This concept can also be applied in contexts other than optical networks. In general, buffering of video frames at an intermediate node, as described above, can take place in any access network having switched video
15 service. Moreover, buffering can also take place at a central office when, for example, the system topology does not include an ON.

C. Channel Surfing and Proactive Streaming

Given the latency that can occur when a user switches among different
20 video transmissions, the practice of scanning multiple transmissions in sequence (analogous to "channel surfing") becomes difficult. This can be addressed by making the switching functionality more intelligent. If, for example, switching is done at the ON (as described above), the ON can be made to sense when channel surfing is taking place.

This is illustrated in FIG. 21. This process begins at step 2105. In step
25 2110, a determination is made as to whether a user has requested some number N of sequential switches within a predetermined window of time. If so, it is assumed that the user is surfing and, in step 2115, the ON proactively sends the next transmission to the user prior to receipt of the actual switch request.
30 The determination of step 2110 can then be repeated, and the next transmission can likewise be sent proactively, etc. When, in step 2110, it is

- 24 -

determined that surfing has stopped (i.e., fewer than N sequential switches within the time window), the next transmission is not sent, and the determination of step 2110 is repeated. This serves to monitor the user for subsequent channel surfing.

5 This concept can also be applied in contexts other than optical networks. In general, detection of sequential switches and anticipation of future switching at an intermediate node, as described above, can take place in any access network having switched video service. Moreover, this functionality can also be placed in a central office when, for example, the
10 system does not include an ON.

B. Other DOCSIS Variations

Other variations on DOCSIS 1.1 can be used for the sake of economy
15 and computational simplicity. In particular, DOCSIS can be implemented without one or more of the features specified by the standard. For example, packet fragmentation/reconstruction and payload header suppression can be omitted, since these functions can be computationally intensive. Likewise, the packet classification function can be limited. These omissions can make
20 processing faster and can in some circumstances increase available bandwidth.

IV. Conclusion

While various embodiments of the present invention have been described above, it should be understood that they have been presented by way
25 of example, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in detail can be made therein without departing from the spirit and scope of the invention. Thus the present invention should not be limited by any of the above-described exemplary
 embodiments.

30

- 25 -

WHAT IS CLAIMED IS:

1. A fiberoptic, ethernet-based time division multiple access (TDMA) communications system, comprising:

a central office;

5 at least one end-user device in optical communication with said central office; and

an aggregating optical node disposed between said central office and said at least one end-user device,

10 wherein a first laser provides optical energy to carry communications between said central office and said optical node, and at least one second laser provides optical energy to carry communications between said optical node and said at least one end-user device.

15 2. The system of claim 1, wherein communications between said optical node and said at least one end-user device comprises a fast ethernet point-to-point connection.

20 3. The system of claim 1, wherein communications between said optical node and said at least one end-user device comprises a shared broadcast connection.

25 4. The system of claim 1, wherein communications between said optical node and said at least one end-user device comprises a gigabit ethernet connection.

5. The system of claim 1, wherein said optical node comprises:
a gigabit ethernet interface to said central office;
a switch, connected to said gigabit ethernet interface, to aggregate traffic from at least one end-user device to said central office;

- 26 -

at least one passive optical network controller, connected to said switch, to arbitrate access of said at least one end-user device to said switch; and

5 at least one physical layer device, to interface at least one end-user device to said at least one passive optical network controller.

6. The system of claim 5, wherein said switch comprises a multiplexer/demultiplexer.

10 7. The system of claim 1, wherein said central office further comprises a physical layer device and a media access controller, wherein said media access controller comprises a gigabit media-independent interface, wherein said gigabit media-independent interface monitors incoming data from said at least one end-user device to detect a field having a known
15 relationship to the start of a frame, for purposes of determining a ranging offset.

8. The system of claim 1,
wherein said optical node comprises a physical layer device
20 and a media access controller, and
wherein said media access controller comprises a gigabit media-independent interface that monitors incoming data from said at least one end-user device to detect a field having a known relationship to the start of a frame, for purposes of determining a ranging offset.

25 9. The system of claim 8, wherein said system is an ethernet-based communications system.

30 10. The system of claim 8, wherein said system is a time division multiple access (TDMA) communications system.

- 27 -

11. The system of claim 1, wherein said central office comprises logic for subscriber service functions, and wherein said optical node comprises logic for bandwidth allocation functions.

5 12. The system of claim 11, wherein said subscriber service functions comprise authentication.

13. The system of claim 11, wherein said subscriber service functions comprise billing functions.

10

14. The system of claim 11, wherein said bandwidth allocation functions comprise timing functions.

15 15. The system of claim 11, wherein said bandwidth allocation functions comprise generation of map messages.

16. The system of claim 1, wherein said at least one end user comprises an adaptive equalizer that cancels noise generated by said first laser.

20 17. The system of claim 16, wherein said noise comprises humming.

18. The system of claim 1, wherein said central office comprises an adaptive equalizer that cancels noise generated by said respective second laser.

25

19. The system of claim 18, wherein said noise comprises humming.

30 20. The system of claim 1, wherein said optical node comprises logic for quality of service functions.

- 28 -

21. The system of claim 20, wherein said quality of service functions comprise traffic management based on prioritization.

5 22. The system of claim 20, wherein said quality of service functions comprise traffic management based on categorization of traffic.

23. The system of claim 1, wherein communications between said central office and said at least one end user device are managed in accordance with Data Over Cable System Interface Specification (DOCSIS) 1.1.

10

24. The system of claim 1, wherein said communications between said central office and said optical node comprise downstream video data, downstream non-video data, and upstream data, and wherein:

15 said downstream digital video is transmitted at a first wavelength;

said downstream non-video data is transmitted at a second wavelength; and

said upstream data is transmitted at a third wavelength, wherein none of said wavelengths are equal.

20

25. The system of claim 1, wherein said communications between said optical node and each of said at least one end-user devices is carried on a different respective frequency band.

26. The system of claim 1, wherein communications to said central office from said at least one end-user device are transmitted on a plurality of subcarriers, wherein each of at least one end-user device transmits said communications to said central office on a different respective one of said subcarriers.

30

- 29 -

27. The system of claim 1, wherein communications from said central office to said at least one end-user device comprise a plurality of channels, communications from said at least one end-user device to said optical node comprise commands to direct one of said channels to one of said at least one end-user, and wherein said optical node comprises logic for directing said one of said channels to said one of said at least one end-user.

28. An access network having switched video service, the network comprising:

10 a central office;

at least one end-user device in communication with said central office; and

an aggregating node disposed between and in communication with said central office and said at least one end-user device;

15 wherein communications from said central office to said at least one end-user device comprise a sequence of frames of video data that are buffered at said node, such that said sequence of frames is accessible by said at least one end-user device.

29. The system of claim 28, wherein said optical node comprises:

logic for detecting repeated access, by one of said at least one end-user devices, of successive channels; and

logic for sending, to said at least end-user device, data transmitted in a next channel, in anticipation of said at least one end-user device attempting to access said next channel.

30. A fiberoptic, ethernet-based time division multiple access communications system, comprising:

a central office;

30 at least one end-user device in optical communication with said central office;

- 30 -

at least one coupler disposed between and in optical communication with said central office and said at least one end-user device; and

5 an aggregating optical node disposed between and in communication with said central office and said at least one end-user device,

wherein downstream communications from said optical node to said at least one end-user device are transmitted on a first wavelength, and upstream communications from said at least one end-user device to said optical node are transmitted on at least one upstream wavelength, wherein
10 each end-user device transmits on its own respective wavelength.

31. The system of claim 30, wherein a first laser provides optical energy to carry communications to said at least one end-user device, wherein said at least one end-user device comprises a first adaptive equalizer that
15 cancels noise generated by said first laser.

32. The system of claim 30, wherein a second laser provides optical energy to carry upstream communications to said central office, wherein said central office comprises an adaptive equalizer that cancels noise
20 generated by said second laser.

33. A fiberoptic access network, comprising:
a central office; and
at least one end-user device in optical communication with said
25 central office,

wherein said communications between said central office and said optical node comprise downstream video data, downstream non-video data, and upstream data, and wherein
said downstream digital video is transmitted at a first wavelength;

30 said downstream non-video data is transmitted at a second wavelength; and

- 31 -

said upstream data is transmitted at a third wavelength,
wherein none of said wavelengths are equal.

34. A fiberoptic access network, comprising:
5 a central office; and
at least one end-user device in optical communication with said
central office,
wherein said communications between said central and each of said at
least one end-user devices is carried on a different respective frequency band.

10 35. A fiberoptic access network, comprising:
a central office; and
at least one end-user device in optical communication with said
central office,
15 wherein communications to said central office from said at
least one end-user device are transmitted on a plurality of subcarriers, wherein
each of at least one end-user device transmits said communications to said
central office on a different respective one of said subcarriers.

20 36. A fiberoptic access network, comprising:
a central office;
at least one end-user device in optical communication with said
central office; and
25 at least one coupler disposed between and in optical
communication with said central office and said at least one end-user device,
wherein downstream communications from said central office to said at least
one end-user device are transmitted on a first wavelength, and upstream
communications from said at least one end-user device to said central office
are transmitted on at least one upstream wavelength, wherein each end-user
30 device transmits on its own respective wavelength.

- 32 -

37. The system of claim 35, wherein a first laser provides optical energy to carry communications, wherein said at least one end user device comprises a first adaptive equalizer that cancels noise generated by said first laser.

5

38. The system of claim 35, wherein a second laser provides optical energy to carry upstream communications, wherein said central office comprises a second adaptive equalizer that cancels noise generated by said second laser.

10

39. A fiberoptic, access communications system, comprising:
a central office; and
at least one end-user device in optical communication with said central office,

15

wherein said central office comprises a physical layer device and a media access controller,

20

wherein said media access controller comprises a gigabit media-independent interface that monitors incoming data from said at least one end-user device to detect a field having a known relationship to the start of a frame, for purposes of determining a ranging offset.

25

40. A fiberoptic, access communications system, comprising:
a central office;
at least one end-user device in optical communication with said central office; and

30

an aggregating optical node disposed between and in communication with said central office,
wherein said optical node comprises a physical layer device and a media access controller, and
wherein said media access controller comprises a gigabit media-independent interface that monitors incoming data from said at least one end-user device to

- 33 -

detect a field having a known relationship to the start of a frame, for purposes of determining a ranging offset.

- 5 41. A fiber optic, ethernet-based time division multiple access (TDMA) communications system, comprising:
- a central office; and
- at least one end-user device in optical communication with said central office,
- 10 wherein communications from said central office to said at least one end-user device are transmitted over a plurality of channels, communications from said at least one end-user device to said central office comprise commands to direct one of said channels to one of said at least one end-user, and
- 15 wherein said central office comprises logic for directing said one of said channels to said one of said at least one end-user.

42. An access network having switched video service, the network comprising:
- a central office; and
- 20 at least one end-user device in communication with said central office;
- wherein communications from said central office to said at least one end-user device comprise a sequence of frames of video data that are buffered at said central office, such that said sequence of frames is accessible
- 25 by said at least one end-user device.

43. An access network having switched video service, the network comprising:
- a central office; and

- 34 -

at least one end-user device in communication with said central office,

wherein said central office comprises:

logic for detecting repeated access, by one of said at least one end-user devices, of successive channels; and

logic for sending, to said at least end-user device, data transmitted in a next channel, in anticipation of said at least one end-user device attempting to access said next channel.

44. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps performed at a headend device:

- a. phase locking a transmit bit rate to a clock of a headend counter; and
- b. transmitting data, at the transmit bit rate, to an end-user device.

45. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps performed at an end-user device:

- a. receiving data at a transmit bit rate that has been phase-locked to a clock of a headend counter; and
- b. deriving a local clock on the basis of the received data.

46. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. receiving transmitted data;
- b. reading a synchronization byte at each of a plurality of points in the transmitted data; and

- 35 -

c. when a threshold number of synchronization patterns have been read successfully in succession, attaining synchronization.

5 47. The method of claim 46, wherein the transmitted data comprises a plurality of MPEG frames, and each of the plurality of points in the transmitted data is the initial byte of one of the MPEG frames.

10 48. The method of claim 46, wherein the transmitted data comprises a plurality of variable-length packets, and each of the plurality of points in the transmitted data is indicated by a pointer in a previous variable-length packet.

15 49. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps, performed at an end-user device:

- a. receiving a grant message from an ethernet switch, wherein the grant message indicates when the end-user device may transmit;
- b. transmitting data at the indicated time; and
- c. withholding transmission of data otherwise.

20

50. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communication system, comprising the following steps:

- 25 a. determining, at an internet protocol layer, an amount of bandwidth to be granted in an unsolicited grant where the amount is an integer multiple of a maximum data unit; and
- b. issuing the unsolicited grant.

30 51. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- 36 -

- a. receiving an unsolicited grant of bandwidth;
- b. determining the amount of granted bandwidth;
- c. determining the amount of granted bandwidth that can be used by whole internet protocol packets;
- 5 d. determining the remaining bandwidth; and
- e. fragmenting a subsequent internet protocol packet such that a fragment of the subsequent internet protocol packet fits the remaining bandwidth.

10 52. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. determining whether the bandwidth needed for non-video data exceeds a default value and whether the demand for video bandwidth is below a threshold value; and
- 15 b. if so, reallocating video bandwidth for non-video data.

20 53. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. receiving a request for bandwidth; and
- b. if the requested bandwidth is available, granting the request.

25 54. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. registering an optical node; and
- b. allocating bandwidth to the optical node based on a fixed assignment policy.
- 30

- 37 -

55. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. registering an optical node;
- 5 b. assigning a fixed amount of bandwidth based on state to the optical node.

56. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. allocating bandwidth to one or more optical nodes; and
- 10 b. dynamically adjusting bandwidth allocations in response to changes in state.

57. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. allocating bandwidth to one or more optical nodes; and
- c. dynamically adjusting bandwidth allocations in
- 20 response to a request, wherein said adjusting is implemented in accordance with an established dynamic service level agreement.

58. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. receiving an allocation of bandwidth; and
- 25 b. sending a packet as to fit the allocation of bandwidth.

59. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps performed at an end-user device:

30

- 38 -

- a. receiving an allocation of bandwidth;
- b. identifying a highest priority packet to be transmitted;
- c. determining whether the highest priority packet fits the allocation of bandwidth;
- 5 d. if so, sending the highest priority packet; and
- e. if not,
 - i) withdrawing the highest priority packet from consideration;
 - 10 ii) identifying the next highest priority packet to be transmitted; and
 - iii) determining whether the next highest priority packet fits the allocation of bandwidth.

15 60. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. sending a request for bandwidth;
- b. receiving a grant; and
- 20 c. sending a packet fitting the grant.

61. The method of claim 60, further comprising the step of:
receiving the grant in an Ethernet frame.

25 62. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, comprising the following steps:

- a. extending DOCSIS to govern transmissions within an optical system; and
- 30 b. allocating bandwidth to one or more optical nodes in accordance with DOCSIS.

- 39 -

- 5
63. The method of claim 62, wherein step (a) comprises:
providing a physical medium responsive to sending
and/or receiving DOCSIS-compliant communications within an optical
system.
64. The method of claim 62, wherein step (a) comprises:
executing subscriber service functions from a hub in
communication with an optical node.
- 10
65. The method of claim 62, wherein step (a) comprises:
operating an optical node to allocate bandwidth.
66. The method of claim 62, wherein step (a) comprises:
operating an optical node to implement timing functions.
- 15
67. The method of claim 62, wherein step (a) comprises:
implementing a request-grant mechanism to allocate
bandwidth.
- 20
68. The method of claim 62, wherein step (a) comprises:
not using fragmentation during transmissions within the optical
system.
- 25
69. The method of claim 62, wherein step (a) comprises:
not using payload header suppression during transmissions
within the optical system.
- 30
70. The method of claim 62, wherein step (a) comprises:
simplifying the classification process during transmissions
within the optical system.

- 40 -

71. The method of claim 62, wherein step (b) is based on an unsolicited grant service.

72. The method of claim 71, wherein step (b) comprises:
5 extending the unsolicited grant service to accept and/or respond to additional requests.

73. A method of optimizing bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system,
10 comprising the following steps performed at an end-user device:

- a. receiving an allocation of bandwidth;
- b. identifying a highest priority packet to be transmitted;
- c. determining whether the highest priority packet fits the allocation of bandwidth;
- 15 d. if so, sending the highest priority packet; and
- e. if not,
 - i) withdrawing the highest priority packet from consideration;
 - ii) identifying the first packet of the next priority level to be transmitted; and
 - 20 iii) determining whether the next highest priority packet fits the allocation of bandwidth.

74. A computer program product comprising computer usable
25 media having computer readable program code means embodied in said media for causing application programs to execute on a computer that improves bandwidth efficiency in a fiberoptic ethernet-based time division multiple access communications system, said computer readable program code means comprising:

30 a first computer readable program code means for causing the computer to determine whether the bandwidth needed for non-video data

- 41 -

exceeds a default value and whether the demand for video bandwidth is below a threshold value; and

a second computer readable program code means for causing the computer to reallocate video bandwidth for non-video data.




FIG. 1




FIG. 2




FIG. 3



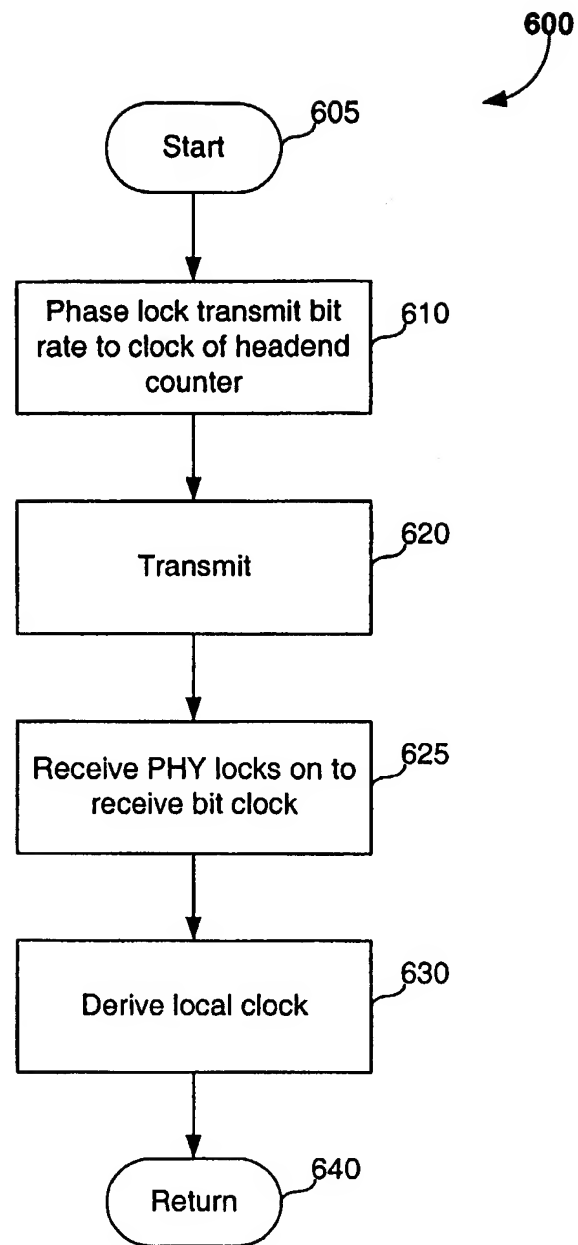


FIG. 4

**FIG. 5**

**FIG. 6**

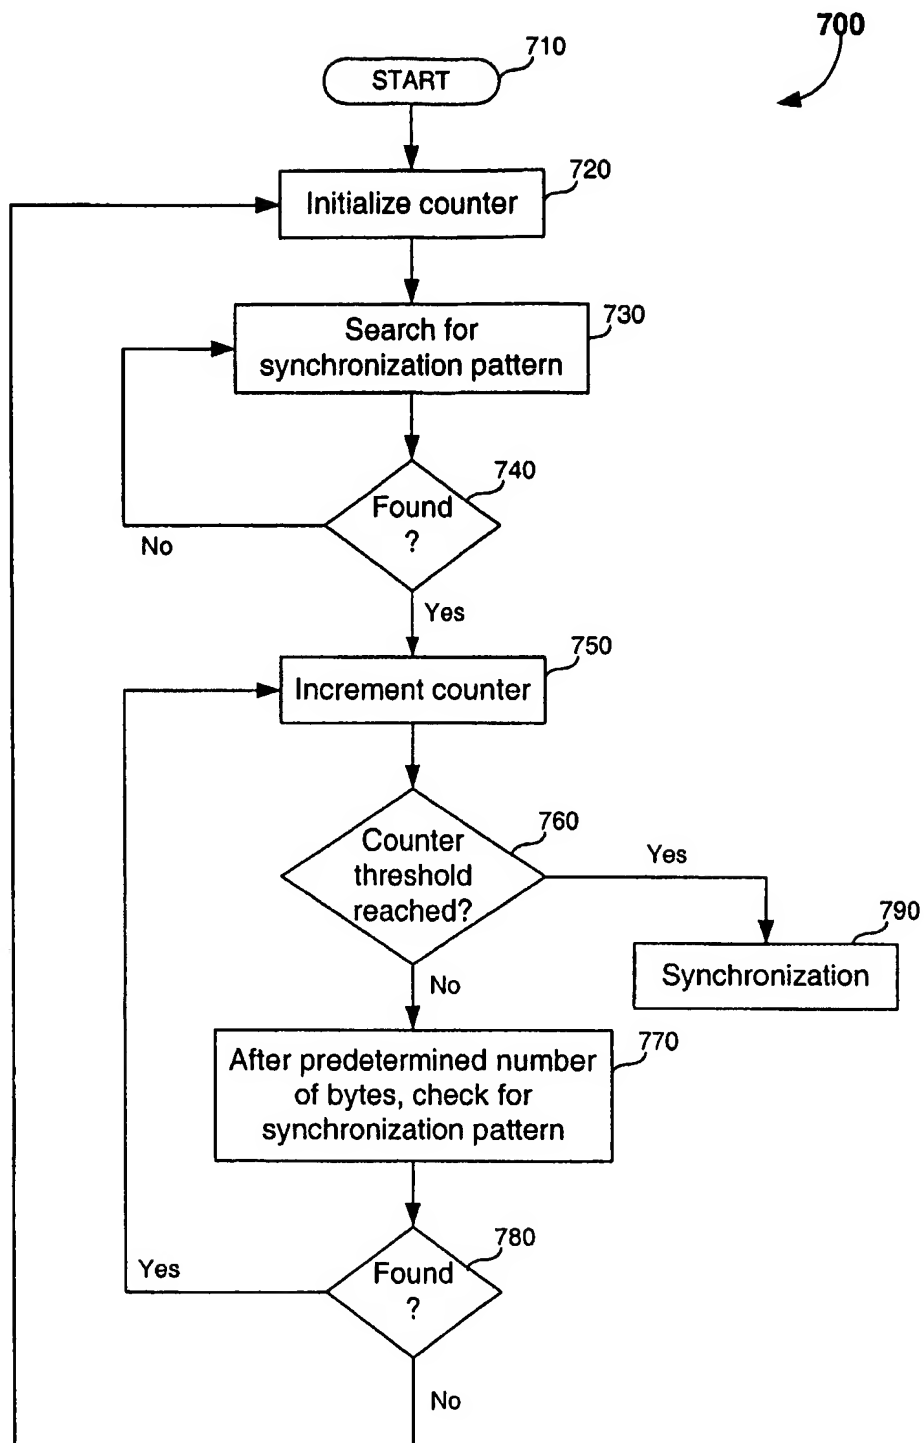


FIG. 7

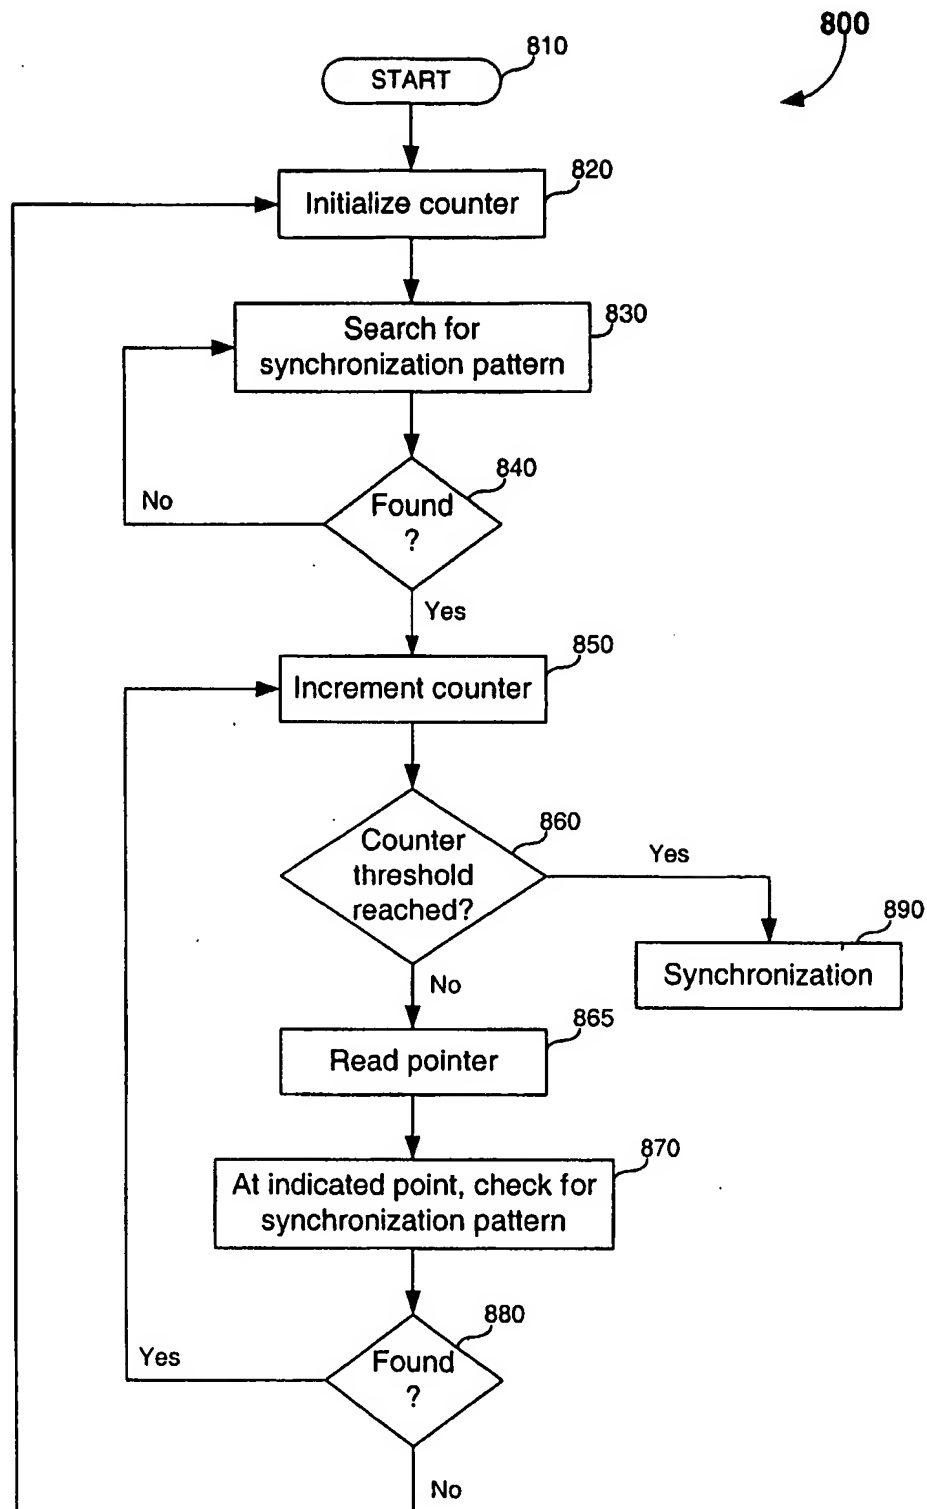


FIG. 8

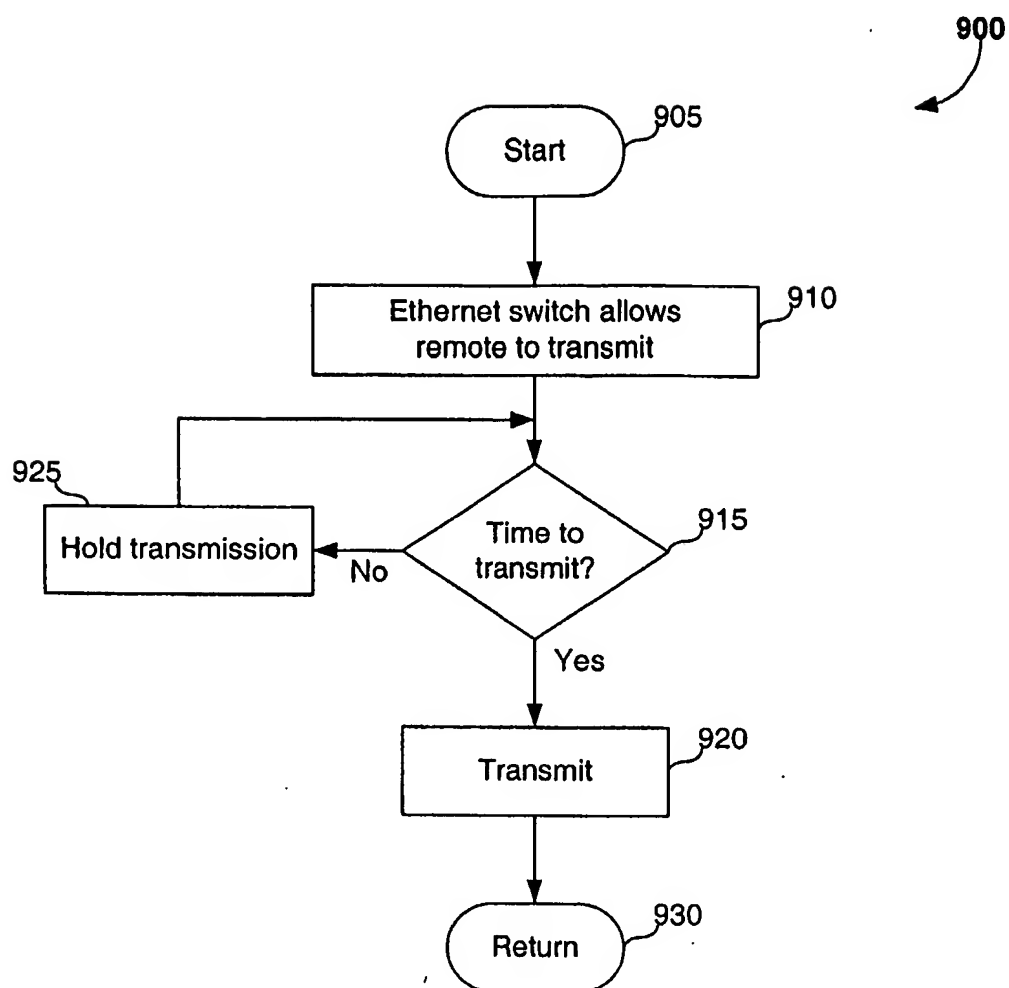




FIG. 9

**FIG. 10A**

**FIG. 10B**




FIG. 11




FIG. 12

1300




FIG. 13




FIG. 14




FIG. 15

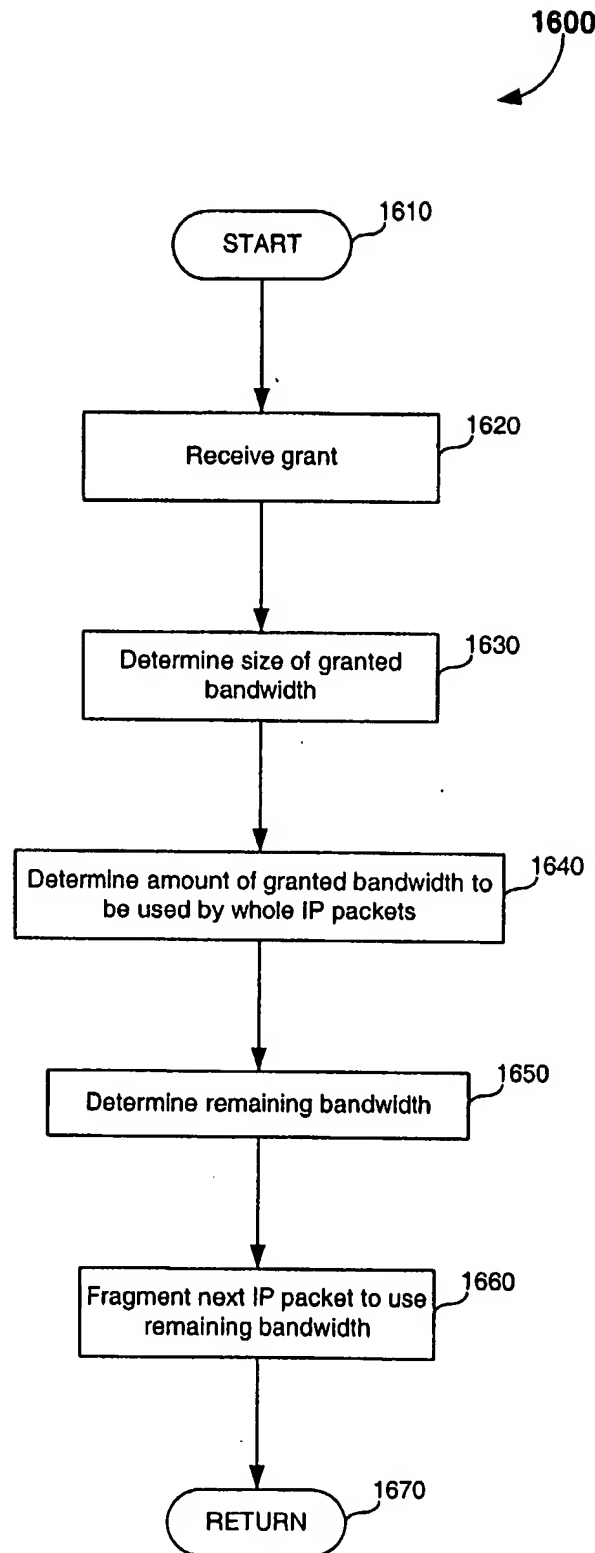


FIG. 16

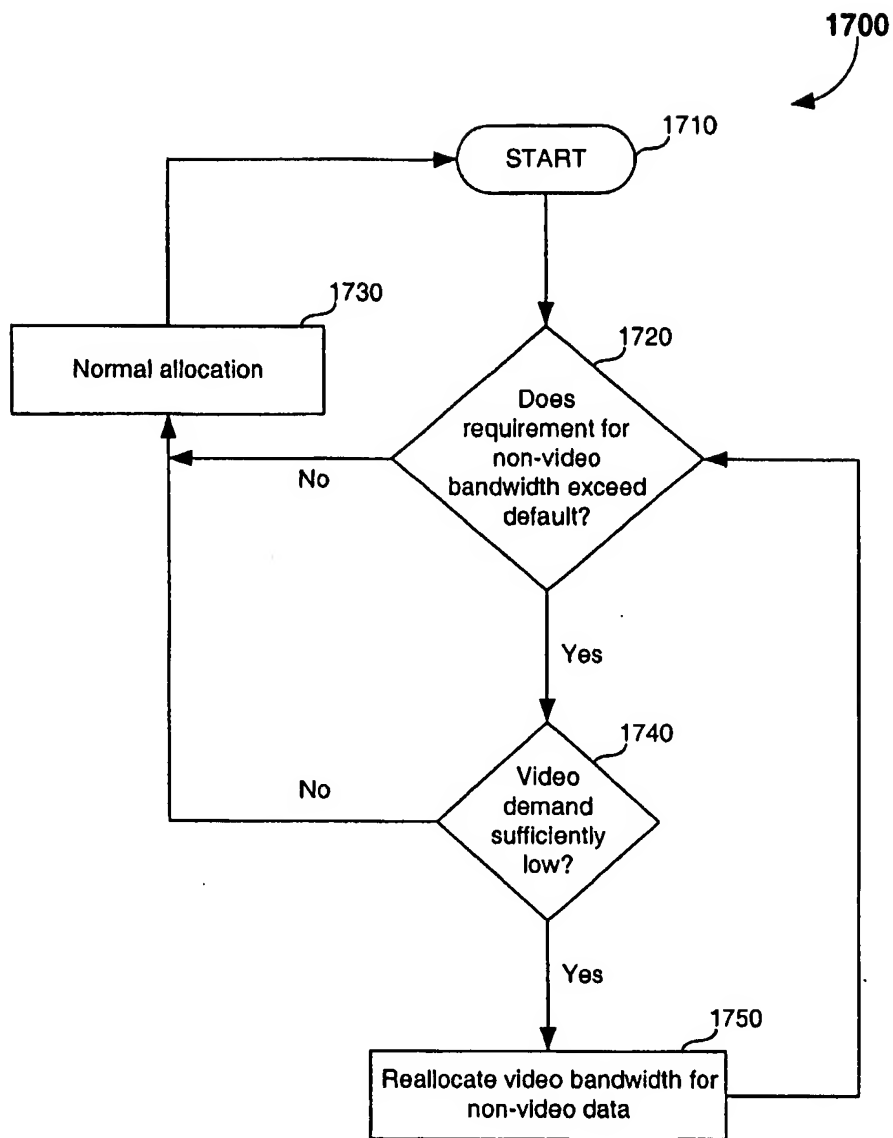


FIG. 17

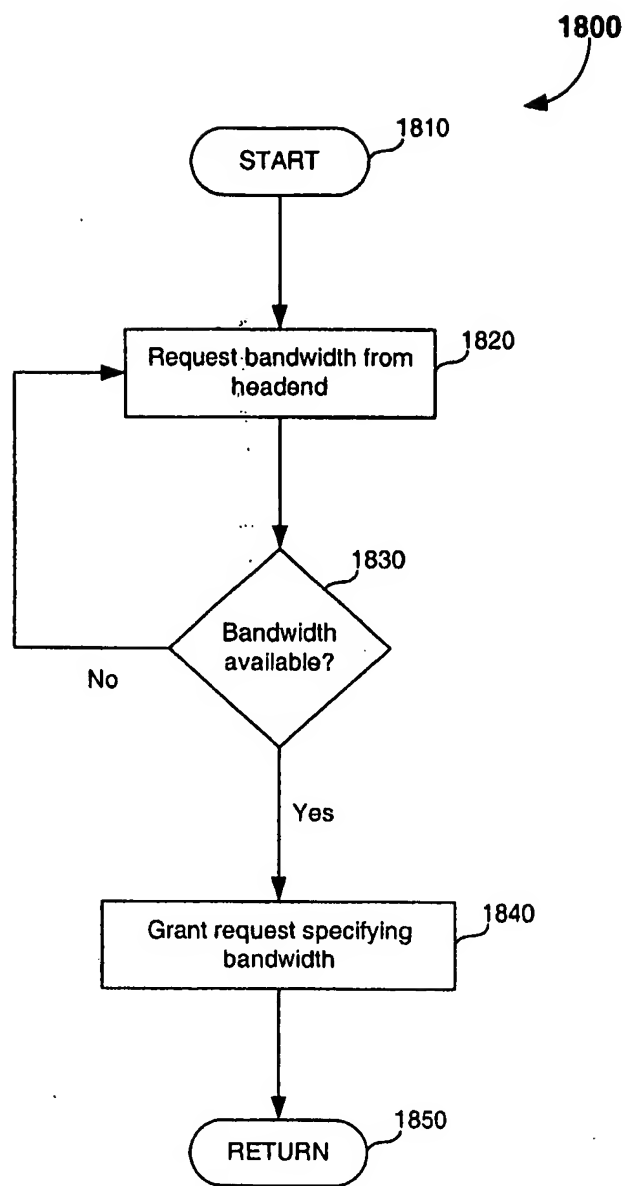


FIG. 18

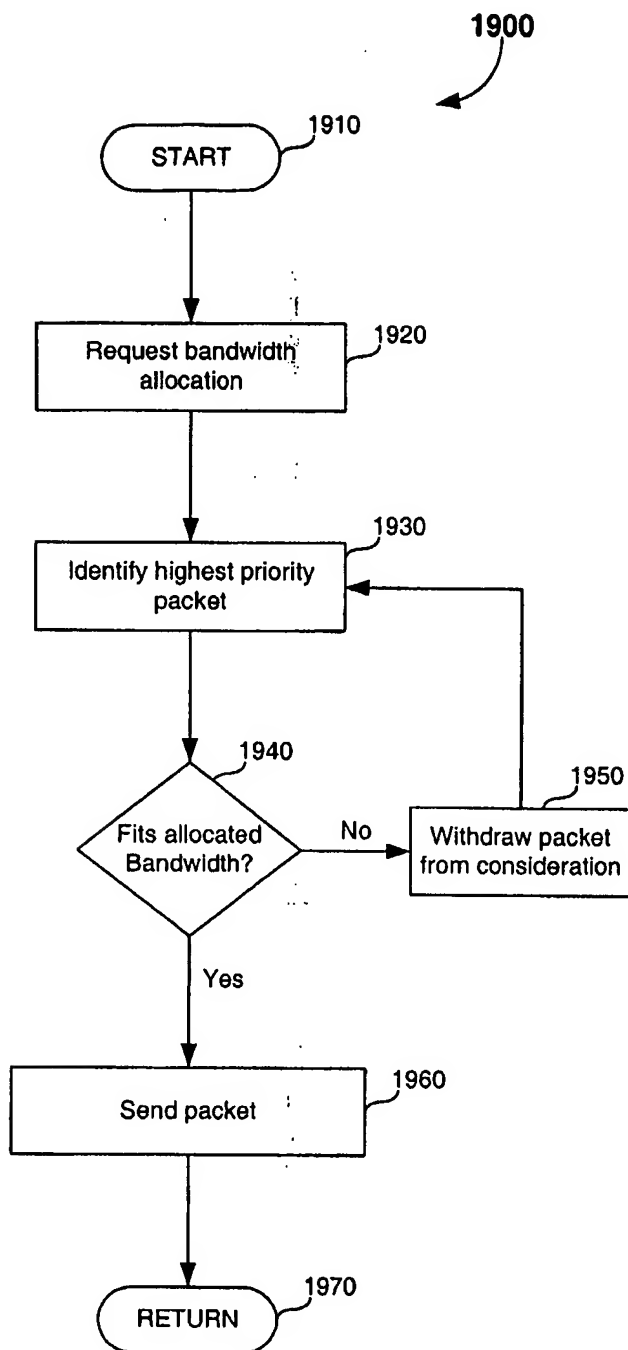


FIG. 19




FIG. 20

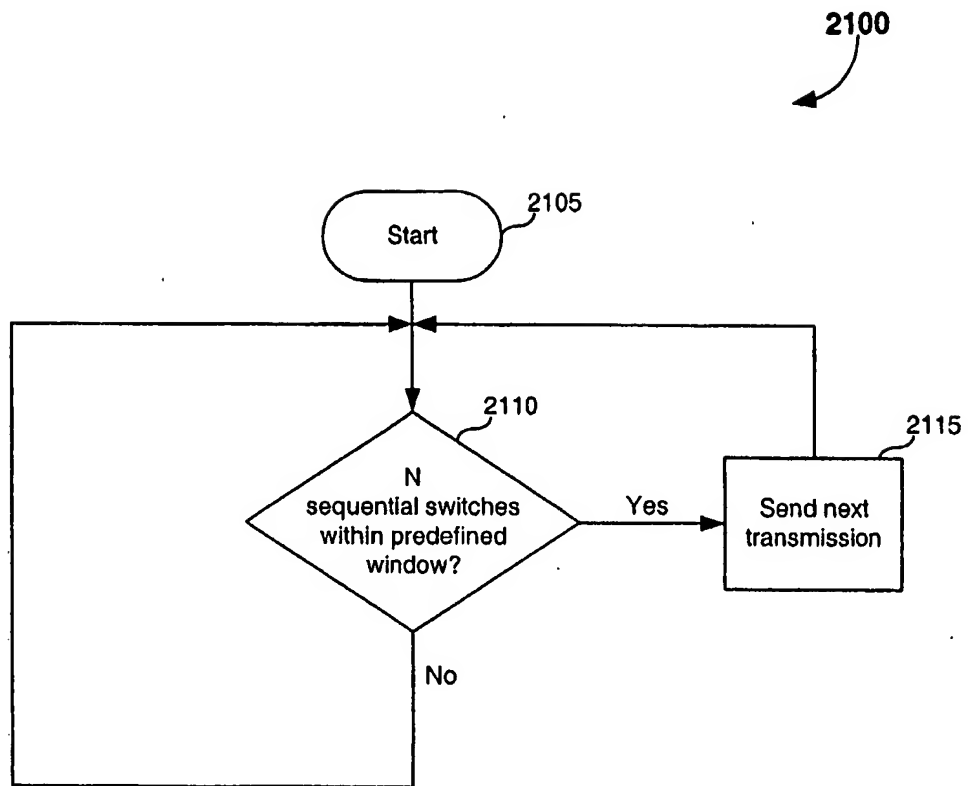


FIG. 21




European Patent Office
Postbus 5818
2280 HV RIJSWIJK
NETHERLANDS
Tel.: +31 70 340 2040
Fax: +31 70 340 3018

Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Grünecker, Kinkeldey,
Stockmair & Schwanhäusser
Anwaltssozietät
Maximilianstrasse 58
80538 München
ALLEMAGNE



EPO Customer Services

Tel.: +31 (0)70 340 45 00

Date

21.09.07

Reference EP30569TE900kap	Application No./Patent No. 04013477.7 - 1237
Applicant/Proprietor SAMSUNG ELECTRONICS CO., LTD.	

COMMUNICATION

The European Patent Office herewith transmits as an enclosure the European search report (under R. 44 or R. 45 EPC) for the above-mentioned European patent application.

If applicable, copies of the documents cited in the European search report are attached.

- ☒ Additional set(s) of copies of the documents cited in the European search report is (are) enclosed as well.

The following specifications given by the applicant have been approved by the Search Division :

- ☒ Abstract ☐ Title
- ☐ The abstract was modified by the Search Division and the definitive text is attached to this communication.

The following figure will be published together with the abstract : 3

Refund of search fee

If applicable under Article 10 Rules relating to fees, a separate communication from the Receiving Section on the refund of the search fee will be sent later.





DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
① X ✓	WO 02/45308 A (ALLOPTIC INC [US]; KUO JERCHEN; GRUIA DUMITRU; KRAMER GLENN; PESAVENTO) 6 June 2002 (2002-06-06) * page 8, line 18 - page 11, line 31 *	1-14	INV. H04Q11/00
② X ✓	WO 03/005156 A (BROADCOM CORP [US]) 16 January 2003 (2003-01-16) * page 14, lines 15-31 * * page 13, lines 23-28; figures 3,4 * * page 19, lines 5-21 * * page 22, lines 10-25 *	1-14	
③ A ✓	US 5 694 234 A (DARCIE THOMAS EDWARD [US] ET AL) 2 December 1997 (1997-12-02) * abstract *	1-14	
			TECHNICAL FIELDS SEARCHED (IPC)
			H04Q H04J
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 7 September 2007	Examiner Meurisse, Wim
CATEGORY OF CITED DOCUMENTS			
<p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 04 01 3477

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07-09-2007

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0245308	A	06-06-2002	CN 1484898 A	24-03-2004
			EP 1342332 A1	10-09-2003
			JP 2004529523 T	24-09-2004
			TW 595153 B	21-06-2004
			US 7031343 B1	18-04-2006
WO 03005156	A	16-01-2003	NONE	
US 5694234	A	02-12-1997	CA 2160823 A1	21-04-1996
			EP 0708540 A1	24-04-1996
			JP 8237223 A	13-09-1996
			US 5680234 A	21-10-1997
			US 5574584 A	12-11-1996